Vote-boosting ensembles

Maryam Sabzevari, Gonzalo Martínez-Muñoz and Alberto Suárez

Abstract—Vote-boosting is a sequential ensemble learning method in which individual classifiers are built on different weighted versions of the training data. To build a new classifier, the weight of each training instance is determined as a function of the disagreement rate of the current ensemble predictions for that particular instance. Experiments using the symmetric beta distribution as the emphasis function and different base learners are used to illustrate the properties and to analyze the performance of these types of ensembles. In classification problems with low or no class-label noise, when simple base learners are used, vote-boosting behaves as if it were an interpolation between bagging and standard boosting (e.g. AdaBoost), depending on the value of the shape parameter of the beta distribution. In terms of predictive accuracy the best results, which are comparable or better than random forests, are obtained with vote-boosting ensembles of random trees.

Index Terms—Ensemble learning, boosting, uncertainty-based emphasis.

arXiv:1606.09458v1 [cs.LG] 30 Jun 2016

1 Introduction

In ensemble learning the outputs of a collection of diverse classifiers are combined to exploit their complementarity, in the expectation that the global ensemble prediction be more accurate than the individual ones [1]. The complementarity of the classifiers is either an indirect consequence of diversity, as in bagging [2] and random forests [3], or can be explicitly favored by design, as in negative correlation learning [4] and boosting [5], [6], [7], [8]. In this manuscript we present vote-boosting, an ensemble learning method of the latter type, in which classifiers progressively focus on different training instances, depending on the degree of uncertainty of their classification. The emphasis is made irrespective of whether the instance is correctly or incorrectly classified. By using the disagreement among the individual predictions to determine the degree of emphasis, one avoids excessive focus on misclassified instances that are outliers, which is one of the reasons for the lack of robustness to class-label noise of standard boosting algorithms.

The article is organized as follows: In Section 2, we provide a review of ensemble methods that are related to the current proposal. Vote-boosting is described in section 3, which includes a derivation of the algorithm as a gradient descent in the functional space of linear combinations of hypothesis. In Section 4, the behavior and performance of vote-boosting ensembles are analyzed using an extensive empirical evaluation on synthetic and real-world classification tasks from different domains of application. Finally, the conclusions of this study are summarized in Section 5.

2 Previous Work

There is a wide variety of methods to build ensembles. Here, we focus on homogeneous ensembles, which are composed of predictors of the same type. Each of the predictors in the ensemble is automatically induced from a training set composed of labeled instances. Finally, a consensus decision is obtained by combining the individual decisions. A wide range of alternatives can be used to carry out this combination [9]. Nevertheless simple strategies, such as averaging for soft outputs or majority voting, if the individual classifiers yield class-labels, are generally effective [10].

Broadly speaking, ensemble methods either explicitly aim to increase the complementarity of the individual classifiers or simply generate diverse predictors using some form of randomization. An example of the former strategy is Negative Correlation Learning [4]. In this method, diversity is favored by simultaneously training all the classifiers in the ensemble: The parameters of the individual classifiers and the weights for the combination of their outputs are determined globally by minimizing a cost function that penalizes coincident predictions. Another example is boosting. Boosting originally refers to the problem of building a strong learner out of a collection of weak learners; i.e. learners whose predictive accuracy is only slightly better than random guessing [5], [7], [8]. AdaBoost is one of the most widely used boosting algorithms [6]. In AdaBoost an ensemble is grown by incorporating classifiers that progressively focus on instances that are misclassified by previous classifiers in the sequence. The individual classifiers are built by applying a learning algorithm that can take into account individual instance weights. The first classifier is obtained by assuming equal weights for all instances. The subsequent classifiers are built using different emphasis on each of the training instances. Specifically, to build the tth classifier in the sequence, the weights of instances that are misclassified by the most recent classifier in the ensemble are increased. Correspondingly, the weights of the correctly classified instances are reduced. The final prediction of the ensemble is determined by weighted majority voting. The weight of an individual classifier in the final ensemble prediction depends on the weighted accuracy of this classifier on the training set. The margin of an instance is defined as the sum of weighted votes for the correct class minus the sum of weighted votes for the most voted incorrect class. Therefore, misclassified instances have negative margins. In AdaBoost, the evolution of the weight of a particular instance is a monotonically decreasing function of its margin [11].

AdaBoost is one of the most effective ensemble methods

[12], [13], [14]. However, it is not robust to class-label noise [15], [16], [17]. Specifically, AdaBoost gives unduly high weights to noisy instances, whose class labels are incorrect. There are numerous studies that address this excessive sensitivity of AdaBoost to class-label noise [7], [18], [19], [20], [21], [22], [23], [24], [25], [26], [27], [28], [29], [30]. A possible strategy is to identify and either remove noisy instances in the training data, or correct their class-labels [18], [19]. Another alternative is to apply explicit or implicit regularization techniques to avoid assigning excessive weight to a reduced group of instances [20], [21], [22], [23]. For instance, the logistic loss function used in LogitBoost [7] gives less emphasis to instances with large negative margins than the exponential loss function used in AdaBoost. In consequence, LogitBoost is generally more robust to class-label noise [31]. In other studies penalty terms are used in the cost function to avoid focusing on outliers or on instances that are difficult to classify [24], [25], [26], [27]. It is possible to also use hybrid weighting methods that modulate the emphasis on instances according to their distance to the decision boundary [32], [33], [34]. Most boosting algorithms use convex loss functions. This has the advantage that the resulting optimization problem can be solved efficiently using, for instance, gradient descent. However, in [35] it is shown that the generalization capacity of boosting variants that use convex loss functions can be severely affected by class-label noise. Alternative non-convex loss functions are used in BrownBoost and other robust boosting variants [28], [29], [30]. In these methods the evolution of the weights is *not* a monotonic function of the margin: instances with small negative margins (i.e. misclassified instances that are close to the decision boundary) are assigned higher weights, as in AdaBoost. However, instances whose margin is negative and large receive lower weights. The rationale for using this type of emphasis is that instances in regions with a large class overlap tend to have small margins. Focusing on these instances is beneficial because the classification boundary can be modeled in more detail. By contrast, large negative margins correspond to misclassified instances that are far from the classification boundary. A robust boosting algorithm should therefore avoid emphasizing these instances, which are likely to be noisy.

Instead of optimization, randomization techniques can be used to generate collections of diverse classifiers. The objective is to build classifiers in such a way that their prediction errors are as independent as possible. If the errors of these classifiers are independent they can be averaged out by the combination process. One example of these types of ensembles is bagging [2]. The individual classifiers in a bagging ensemble are built by applying a fixed learning algorithm to independent bootstrap samples drawn from the original training data. In class-switching ensembles, each member is built using a perturbed versions of the original training set, in which where the class labels of a fraction of instances are modified at random [36]. Alternatively, diverse classifiers can be built by including some randomized step (or steps) in the learning algorithm itself. For instance, in one of the earliest works on ensembles of neural networks [37] one takes advantage of the presence of multiple local minima in the optimization process that is used to tune the synaptic weights and averages over the predictions of the neural networks that result from different initializations. Random forests [3], which are one of the most effective ensemble methods [14], are built using a combination of data randomization and randomization in the learning algorithm: The ensemble classifiers are random trees, which are trained on bootstrap replicates of the original training dataset. Other effective classifiers of this type are rotation forests [38] and ensembles of extremely randomized trees [39].

The aggregation of predictions by simple majority voting in parallel randomized ensembles can be analyzed in detail [37], [40], [41], [42]. An important conclusion of this analysis is that instances for which the disagreement of the individual predictions is large are more difficult to classify. They not only exhibit higher error rates but also require querying more classifiers to reach a stable ensemble prediction. This observation suggests the design of a boosting algorithm in which more emphasis is given to uncertain instances, irrespective of whether they are correctly or incorrectly classified. Using this type of emphasis one naturally avoids focusing on incorrectly classified instances for which most of the individual predictions coincide, but are far from the decision boundary. Using these ideas, vote-boosting, which is described in detail in the following section, assigns weights to the training instances as a function of the disagreement of the ensemble base classifiers votes.

3 VOTE-BOOSTING

Consider the problem of automatic induction of a classification system from labeled training data. The original training set is composed of attribute class-label pairs $\mathcal{D}_{train} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N_{train}}$, where $\mathbf{x} \in \mathcal{X}$. In this article we focus on binary classification tasks, in which $y \in \{-1, 1\}$. Problems with multiple classes can be addressed with any of the strategies that are used in combination with AdaBoost for this purpose [8].

Let $\{f_\tau(\cdot)\}_{\tau=1}^{t}$ be a partially-grown ensemble of size t. The τth classifier in the ensemble is a function $f_\tau : \mathbb{R}^D \rightarrow \{-1, 1\}$ that maps a vector of attributes $\mathbf{x}$ to a class label $f_\tau(\mathbf{x})$. This function is obtained by applying a base learning algorithm to a training set, taking into account the individual instance weights $\left\{w_i^{[\tau]}\right\}_{i=1}^{N_{train}}$.

The predictions of the individual classifiers can be aggregated by weighted averaging to obtain the prediction of the ensemble

$$F_t(\mathbf{x}) = \sum_{\tau=1}^{t} \alpha_\tau^{[t]} f_\tau(\mathbf{x}), \quad F_t \in \{-1, 1\}, \tag{1}$$

where $\alpha_\tau^{[t]} \geq 0$ is the weight of the individual prediction of the τ classifier. These weights are normalized $\sum_{\tau=1}^{t} \alpha_\tau = 1$. In (unweighted) majority voting one assumes that all the predictions have the same weight $\alpha_\tau^{[t]} = 1/t$. Based on this aggregated output, the final prediction of the ensemble of size t is

$$H_t(\mathbf{x}) = sign\left(F_t(\mathbf{x})\right). \tag{2}$$

Assuming that t is odd, $H_t(\mathbf{x}) \in \{-1, 1\}$. At this stage of the ensemble construction process, instance $\mathbf{x}$ can be characterized by $t_+(\mathbf{x})$ and $t_-(\mathbf{x}) = t - t_+(\mathbf{x})$, the counts

of positive and negative votes, respectively. The fractions of votes in each class are

$$\pi_{\pm}^{[t]}(\mathbf{x}) = \frac{t_{\pm}(\mathbf{x})}{t}; \quad \pi_{+}^{[t]}(\mathbf{x}) + \pi_{-}^{[t]}(\mathbf{x}) = 1. \tag{3}$$

These values can be used to quantify the level of certainty of the ensemble prediction. Values $\pi_{+}^{[t]}(\mathbf{x})$ close to 0 or 1 correspond to instances for which the predictions of most ensemble classifiers coincide. Instances whose classification by the ensemble is uncertain are characterized by $\pi_{+}^{[t]}(\mathbf{x})$ close to $1/2$. In vote-boosting the learning process progressively focuses precisely on these uncertain instances, which are close to the classification boundary. To this end, a new ensemble classifier is built on a training set in which such instances are given more weight. In contrast to standard boosting algorithms, the degree of emphasis depends on the disagreement among the predictions of the individual classifiers not on whether these predictions are erroneous. The advantage of this type of emphasis is that it does not consider the class label of the instance. In this manner one avoids focusing on instances on which most classifiers make coincident predictions that are incorrect. These instances are likely to be outliers. Giving them excessive weight tends to misguide the learning process and typically leads to a deterioration of the generalization performance of the ensemble. Furthermore, an emphasis based on votes can be used to boost base learners that attain low errors on the training set, such as unpruned decision trees.

The pseudo-code of the proposed vote-boosting algorithm is presented in Fig. 1.

Algorithm 1: Vote-boosting algorithm with resampling

Input:
$\mathcal{D}_{train} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N_{train}},\ \mathbf{x}_i \in \mathcal{X},\ y_i \in \{-1, 1\}$
T % Ensemble size
$\mathcal{L}$ % Base learning algorithm
$g(p)$ % Non-negative emphasis function ($0 \leq p \leq 1$)

$F_0(\cdot) \leftarrow 0$
$t_+(\mathbf{x}_i) \leftarrow 0\ \forall i = 1, ..., N_{train}$
$w_i^{[1]} \leftarrow \frac{1}{N_{train}}\ \forall i = 1, ..., N_{train}$
for $t \leftarrow 1$ ***to*** T **do**
  $f_t(\cdot) \leftarrow \mathcal{L}\left(\mathcal{D}_{train}, \mathbf{w}^{[t]}\right)$
  $t_+(\mathbf{x}_i) \leftarrow t_+(\mathbf{x}_i) + \mathbb{I}(f_t(\mathbf{x}_i) > 0)\ \forall i = 1, ..., N_{train}$
  $\pi_+(\mathbf{x}_i) = \frac{t_+(\mathbf{x}_i)+1}{t+2}\ \forall i = 1, ..., N_{train}$
  $w_i^{[t+1]} \leftarrow g\left(\pi_+(\mathbf{x}_i)\right)\ \forall i = 1, ..., N_{train}$
  Normalize $\mathbf{w}^{[t]}$
  $F_t(\cdot) \leftarrow F_{t-1}(\cdot) + f_t(\cdot)$

Output: $H_T(\cdot) = sign(F_T(\cdot))$

The final ensemble is composed of T classifiers, each of which is built by applying the base learning algorithm $\mathcal{L}$ on the training set with different sets of instance weights. The weights of the instances can be taken into account using weighted resampling

$$\begin{aligned} \mathcal{D}_{train}^{[t]} &= \text{Sample}\left(\mathcal{D}_{train}, \mathbf{w}^{[t]}\right) \\ f_t(\cdot) &\leftarrow \mathcal{L}\left(\mathcal{D}_{train}^{[t]}\right). \end{aligned} \tag{4}$$

For the induction of the first ensemble classifier all instances are assigned the same weight. These weights are updated at each iteration according to the tally of votes: Assuming that instance $\mathbf{x}_i$ has received $t_+(\mathbf{x}_i)$ votes for the positive class at the tth iteration, we use the Laplace estimator of the probability that a classifier in the ensemble outputs a particular class prediction

$$\pi_{\pm}(\mathbf{x}_i) = \frac{t_{\pm}(\mathbf{x}_i) + 1}{t + 2}. \tag{5}$$

Finally the weights are updated according to

$$w_i^{[t+1]} = \frac{1}{\mathcal{Z}_{t+1}} g\left(\pi_+(\mathbf{x}_i)\right), \quad \text{for } i = 1, ..., N_{train}, \tag{6}$$

where $g : [0, 1] \rightarrow \mathbb{R}^+$ is an emphasis function, which is non-negative, and

$$\mathcal{Z}_{t+1} = \sum_{i=1}^{N_{train}} g\left(\pi_+(\mathbf{x}_i)\right) \tag{7}$$

is a normalization constant. These weights are then used to build the next classifier in the ensemble.

A natural choice for the emphasis function is the probability density of the beta distribution with shape parameters a, b

$$g(p) \equiv \beta\left(p; a, b\right) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} p^{a-1}(1-p)^{b-1}, \quad 0 \leq p \leq 1. \tag{8}$$

For this particular choice of emphasis function, the weights at the $(t + 1)$ iteration are updated according to

$$w_i^{[t+1]} = \frac{1}{Z_{t+1}} \beta\left(\pi_+(\mathbf{x}_i); a, b\right) \tag{9}$$

If the class distributions are not strongly imbalanced, the choice $a = b$, in which the two classes are handled in a symmetrical manner, is generally appropriate. In Fig. 1 the density profiles of the symmetric beta distribution for different values of a=$b \in \{0.25, 0.75, 1, 1.5, 2.5, 5, 10, 20, 40\}$ are shown. If $a = b = 1$, the beta distribution is the uniform distribution (plot in the first row, third column of Fig. 1). If $a = b < 1$ values close to 0 and 1 have higher probability density. For $a = b > 1$ the distribution becomes unimodal, with a maximum at 0.5. In this range, the higher the values of $a = b$, the more concentrated becomes the probability around the mode. In consequence, when values of a and b above 1 are used, vote-boosting emphasizes uncertain instances and reduces the importance of those instances on which most classifiers agree.

Note that the Laplace correction in (5) avoids evaluations of the beta distribution at the boundaries of its support: It is common, especially in the first iterations of the algorithm, when the ensemble is still small, that all the classifiers predict the same class label for some instances. In such cases, the fraction of positive votes is either 0 or 1. Except for $a = b = 1$, the value of the beta distribution at these points is either zero or infinity. In the case of zero density values, those instances would be assigned zero weight in the next iteration of the algorithm and would be effectively removed from the sample. The opposite effect would be the result of some instances being assigned infinite weights. To

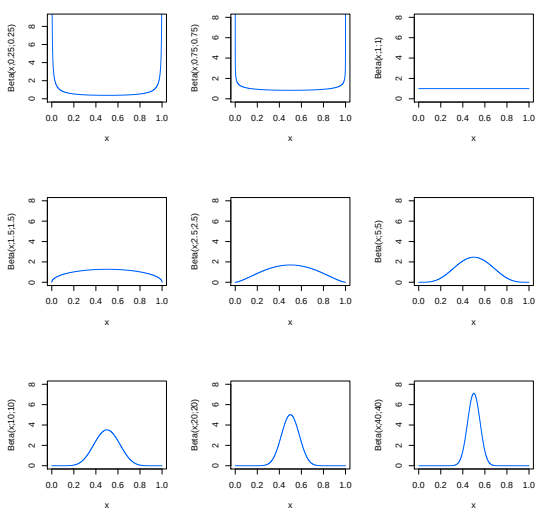

Fig. 1. Symmetric beta distribution with α=β=[0.25,0.75,1,1.5,2.5,5,10,20,40]

avoid these difficulties, the Laplace correction has been used in the estimation of class prediction probabilities (5).

If the emphasis function is the beta distribution with parameters $a = b = 1$, all instances are given the same importance and the proposed algorithm is equivalent to bagging [2] if weighted resampling is used. The emphasis on uncertain instances that one obtains when higher values of a and b are used is similar to the error-based emphasis given by AdaBoost. The reason for this is that uncertain instances are generally more difficult to classify and, in consequence, are more likely to be incorrectly classified. In the regime $a = b < 1.0$ one emphasizes instances that are far from the classification boundary, which, as will be illustrated in the section on experiments, is not in general an effective emphasizing strategy but can provide some regularization in complex problems.

3.1 A functional gradient interpretation of vote-boosting

Similarly to other boosting methods, vote-boosting can be viewed as a gradient descent algorithm in the hypothesis space of linear combinations of predictors [43]. Consider an ensemble of t predictors $\{f_\tau(\cdot)\}_{\tau=1}^{t}$. The global ensemble prediction on instance $\mathbf{x}$ is of the form

$$H_t(\mathbf{x}) = sign[F_t(\mathbf{x})], \tag{10}$$

where

$$F_t(\mathbf{x}) = \frac{1}{t}\sum_{\tau=1}^{t} f_\tau(\mathbf{x}), \quad F_t(\mathbf{x}) \in [-1,1] \tag{11}$$

The fraction of votes for the positive class can be expressed in terms of this quantity

$$\pi_+^{[t]}\mathbf{x}) = \frac{1+F_t(\mathbf{x})}{2}. \tag{12}$$

Consider the cost functional

$$C[F] = \frac{1}{N_{train}}\sum_{i=1}^{N_{train}} y_i c(F(\mathbf{x}_i)), \tag{13}$$

where $y_i c(F(\mathbf{x}_i)$ is the contribution to the cost associated to instance $\mathbf{x}_i$.

For a given F_t, our goal is to find the function $f \in \mathcal{F}$ that minimizes $C[F_{t+1}]$, where

$$F_{t+1}(\mathbf{x}) = F_t(\mathbf{x}) + \frac{1}{t+1}\left(f(\mathbf{x}) - F_t(\mathbf{x})\right). \tag{14}$$

Assuming the change in the value of the cost functional when the ensemble incorporates the new classifier f is small

$$\delta C[F] \equiv C\left[F_{t+1}(\mathbf{x})\right] - C\left[F_t(\mathbf{x})\right] \propto \left[\sum_{i=1}^{N_{train}} y_i c'(F_t(\mathbf{x}_i)) f(\mathbf{x}_i) - \sum_{i=1}^{N_{train}} y_i c'(F_t(\mathbf{x}_i)) F_t(\mathbf{x}_i)\right]. \tag{15}$$

The second term in this expression does not depend on f. Therefore, to lowest order, minimizing the cost functional is equivalent to minimizing

$$\begin{aligned}
&\sum_{i=1}^{N_{train}} y_i c'(F_t(\mathbf{x}_i)) f(\mathbf{x}_i) \\
&\quad = \sum_{i:y_i=f(x_i)} c'(F_t(\mathbf{x}_i)) - \sum_{i:y_i\neq f(x_i)} c'(F_t(\mathbf{x}_i)) \\
&\quad = \sum_{i=1}^{N_{train}} c'(F_t(\mathbf{x}_i)) - 2\sum_{i:y_i\neq f(x_i)} c'(F_t(\mathbf{x}_i)) \\
&\quad = 2\sum_{i=1}^{N_{train}} c'(F_t(\mathbf{x}_i)\left(\frac{1}{2} - \sum_{i:y_i\neq f(x_i)} w_i^{[t+1]}\right),
\end{aligned}$$

where

$$w_i^{[t+1]} = \frac{1}{\sum_{j=1}^{N_{train}} c'\left(F_t(\mathbf{x}_j)\right)} c'\left(F_t(\mathbf{x}_i)\right). \tag{16}$$

Assuming that $c(x)$ is a monotonic non-increasing function, then $-c'(x)$ is non-negative and the values $\left\{w_i^{[t+1]}\right\}_{i=1}^{N_{train}}$ can be thought of as a set of instance weights. The denominator in (16) ensures that these weights are normalized

$$\sum_{i=1}^{N_{train}} w_i^{[t+1]} = 1. \tag{17}$$

Under these conditions the $(t+1)$th predictor in the ensemble is the minimizer of the *weighted* training error

$$f_{t+1} = \arg\min_{f\in\mathcal{F}} \sum_{i:y_i\neq f(x_i)} w_i^{[t+1]}. \tag{18}$$

where $\mathcal{F}$ is the functional space of the base learners.

The conditions necessary for the interpretation of (16) as weights are fulfilled if the cost function is expressed in terms of a cumulative distribution function $G(p)$ defined in $p \in [0,1]$

$$c\left(F(\mathbf{x})\right) = 2\left[G(1/2) - G\left(\frac{1+F(\mathbf{x})}{2}\right)\right]. \tag{19}$$

Because of the monotonicity of $G(p)$, when the prediction $F(\mathbf{x}_i)$ is incorrect (i.e. $y_i F(\mathbf{x}_i) < 0$) the contribution to the cost functional $y_i c\left(F(\mathbf{x}_i)\right)$ is positive. Assuming this form for $c\left(F(\mathbf{x})\right)$, its derivative is

$$c'\left(F(\mathbf{x})\right) = -g\left(\frac{1+F(\mathbf{x})}{2}\right), \tag{20}$$

where $g(p) = G'(p)$ is the corresponding probability density, which is non-negative. Therefore, the weights of the training instances are

$$w^{[t+1]} = \frac{g\left(\pi_t(\mathbf{x}_i)\right)}{\sum_{j=1}^{N_{train}} g\left(\pi_t(\mathbf{x}_j)\right)}, \quad i = 1, 2, \ldots \tag{21}$$

where the density $g(p)$ plays the role of an emphasis function. Note that $\delta C[F] < 0$ only if the weighted training error is lower than $1/2$. In vote-boosting the emphasis is made on instances close to the classification boundary, which, in binary classification problems, is defined as the set of points in which the disagreement among the individual predictions is largest (i.e. half of the classifiers predict one class and half of the classifiers the other one). In the problems investigated, this type of emphasis does not lead to an increase of the weighted training error above this limit. In fact the predictive accuracy of the ensemble generally increases with its size. Therefore, in the implementation of vote-boosting, tested in the experiments presented in the next section, the ensembles are built to a fixed size, which is sufficiently large for convergence to the asymptotic error level [41].

4 Experimental Results

In this section we present the results of an empirical analysis of vote-boosting ensembles. Different sets of experiments have been carried out to characterize the behavior of these types of ensembles and to measure their accuracy in a wide range of classification tasks from different areas of application. In all experiments the symmetric beta distribution has been used as the emphasis function. The relative importance of the training instances in the induction of an individual ensemble classifier is controlled by $a = b$, the shape parameter of the symmetric beta distribution.

A first set of experiments is carried out to investigate the relationship of vote-boosting with bagging and AdaBoost. These experiments show that, in some cases, vote-boosting can be seen as an interpolation between bagging ($a = b = 1.0$) and AdaBoost (for higher values of $a = b$).

In a second set of experiments the accuracy of vote-boosting ensembles is compared with bagging, AdaBoost and random forest. From the results obtained in this benchmarking exercise we conclude that vote-boosting ensembles composed of random trees achieve state-of-the-art classification accuracy rates, which are comparable or superior to random forest and AdaBoost in the problems investigated [14].

4.1 Vote-boosting as an interpolation between bagging and AdaBoost

The objective of the experiments presented in this subsection is to investigate how the characteristics of vote-boosting ensembles change when different levels of emphasis on the uncertain training instances are considered. As discussed earlier, when uniform emphasis is made, vote-boosting is equivalent to bagging. In most of the problems analyzed, when simple or regularized learners (decision stumps, pruned CART trees) are used to build the ensemble, stronger emphasis on uncertain instances results in a behavior that is similar to AdaBoost. In fact, for these problems, vote-boosting performs as an interpolation between bagging and AdaBoost depending on the strength of the emphasis on uncertain instances. Complex unregularized classifiers such as unpruned CART trees or random trees cannot be used in combination with AdaBoost because they achieve training error rates that are either zero or very low. In consequence, the weights for the instances in the next iteration of AdaBoost would be ill defined. Vote-boosting does not have this limitation and, as shown in the following subsection, can be used to build ensembles of these types of classifiers with good predictive accuracy.

To investigate the relations between vote-boosting and AdaBoost we first present the results of an experiment in the binary classification problem *Twonorm* using decision stumps as base learners. In *Twonorm* instances are drawn from two unit-variance Gaussians in 20 dimensions whose means are $(a, a, ..a)$ for one class and $(-a, -a, .. - a)$ for the other one, with $a = 2/\sqrt{(}20)$ [44]. Note that this is not a trivial task for decision stumps because, as individual classifiers, they can model only class boundaries that are parallel to the axes. In the experiments performed, the training set is composed of 500 independently generated examples. Different vote-boosting ensembles composed of 100 stumps were built using the symmetric beta distribution for emphasis, with $a = b \in \{1.0, 2.0, 30.0\}$. An AdaBoost ensemble composed of 100 stumps was also built using the same training data. The final weights given to the training instances in the different ensembles were recorded and subsequently ranked. Ties were resolved by randomizing the corresponding ranks. In Fig. 2 we display the scatter plot of these ranks for AdaBoost on the horizontal axis and for vote-boosting on the vertical axis with $a = b = 1.0$ in (a), $a = b = 2.0$ in (b), and $a = b = 30.0$ in (c).

When $a = b = 1.0$ is used, all instances are assigned the same weight and no correlations between the weight ranks given by vote-boosting and AdaBoost can be observed. Therefore, the points corresponding to the training instances appear uniformly distributed in Fig. 2 (a). As the values of $a = b$ are increased, points tend to cluster around the diagonal. This is a consequence of the fact that the ranks of the weights given by both types of ensembles become similar. Comparing Figs. 2 (a), (b) and (c) it is apparent that the correlations between the weight ranks become stronger as $a = b$ increases. In particular for $a = b = 30.0$ vote-boosting and AdaBoost give similar emphasis, in spite of the fact that the former does not make use of class labels to decide whether an instance should be given more weight, whereas the latter does. The reason for this coincident emphasis is that, in this simple problem, the ensemble classifiers are more likely to disagree precisely in the instances that are incorrectly classified.

However, if class-label noise is injected in the problem, the weighting schemes of AdaBoost and vote-boosting become different: Vote-boosting maintains the focus on instances in the boundary region, in which classes overlap and the disagreement rates among the ensemble preditions are highest. By contrast, AdaBoost tends to give more weight to those instances whose class label has been modified. Focusing on these noisy instances is misleading and eventually impairs the generalization capacity of AdaBoost. To illus-

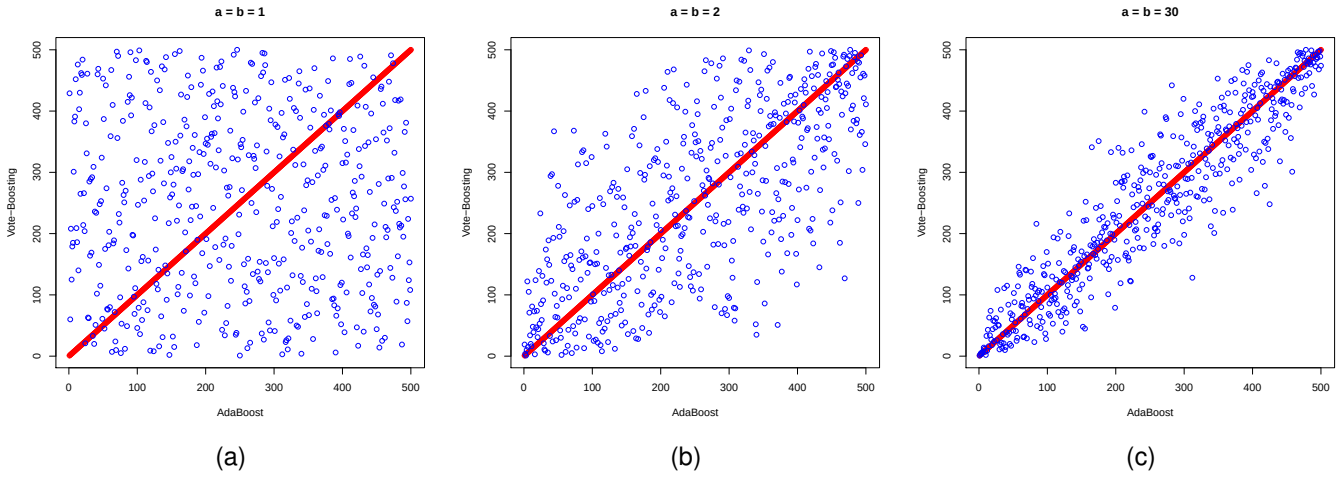

Fig. 2. Weight ranks of the training instances for vote-boosting and AdaBoost in *Twonorm* (a) $a = b = 1.0$, (b) $a = b = 2.0$, (d) $a = b = 30.0$

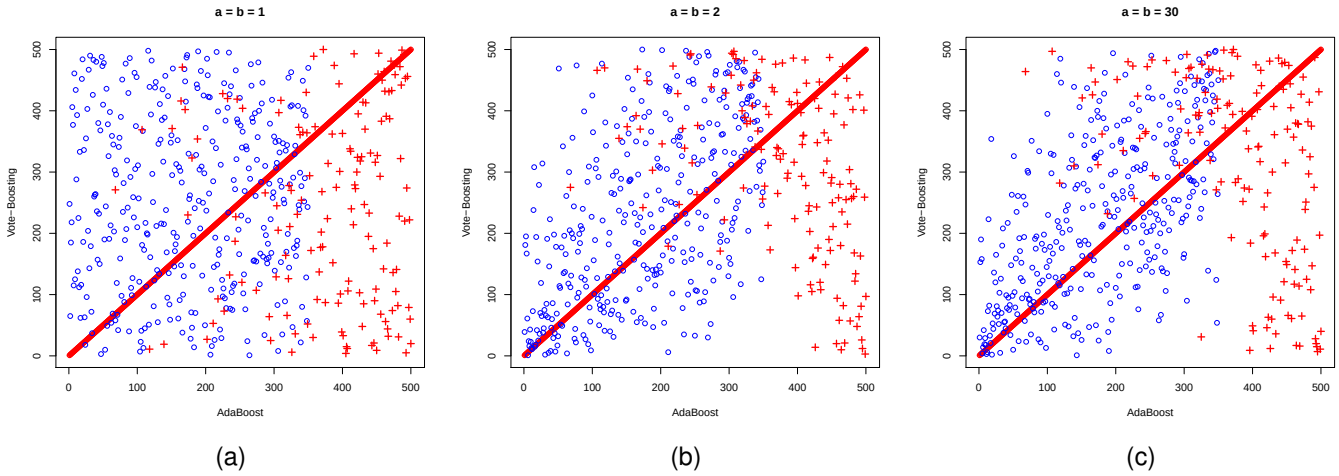

Fig. 3. Weight ranks of the training instances for vote-boosting and AdaBoost in *Twonorm* with 30% class-label noise, (a) $a = b = 1.0$, (b) $a = b = 2.0$, (c) $a = b = 30.0$

trate this observation the experiment was repeated injecting class-label noise in the training data. Specifically, 30% of the training examples were selected at random and their class labels flipped, as in the noisy completely at random model described in [45]. The results of these experiments are shown in Fig. 3. Instances whose class label has been switched are marked with a red cross in these plots. For $a = b = 1.0$, no correlation is observed between the ranks of the weights given by vote-boosting and AdaBoost. However, instances whose class label has been switched, which are distributed uniformly in the vertical direction, tend to appear on the right-hand side of the plots. This means that they receive special emphasis in Adaboost but not in vote-boosting. Increasing the value of $a = b$ pushes the unperturbed instances towards the diagonal but not the perturbed ones. However, the correlation between the ranks is less marked than in the noiseless case, because of the interference of the noisy instances.

A second batch of experiments was carried out to analyze the behavior of vote-boosting as a function of the strength of the emphasis that is applied to uncertain instances. Using the symmetric beta distribution as emphasis function, we analyze how the learning curves, which trace the dependence of the error as a function of the size of the ensemble, depend on the value of the shape parameter. The values explored are $a = b \in \{0.25, 0.5, 0.75, 1.0, 1.25, 1.5, 2.5, 5, 10, 20, 40\}$. The experiments were made on the classification tasks *Twonorm* and *Pima*. These tasks have been chosen because the different prediction accuracies of bagging and AdaBoost. In *Twonorm*, AdaBoost significantly outperforms bagging. By contrast in *Pima*, which is a very noisy task, bagging is more accurate than AdaBoost.

Figures 4 and 5 display the learning curves for bagging, AdaBoost and vote-boosting using different values of the shape parameter, in the classification tasks *Twonorm* and *Pima*, respectively. The plots on the left-hand side show the results for the training error. The plots on the right-hand side correspond to the test error curves.

When $a = b = 1.0$, all instances are given equal weights.

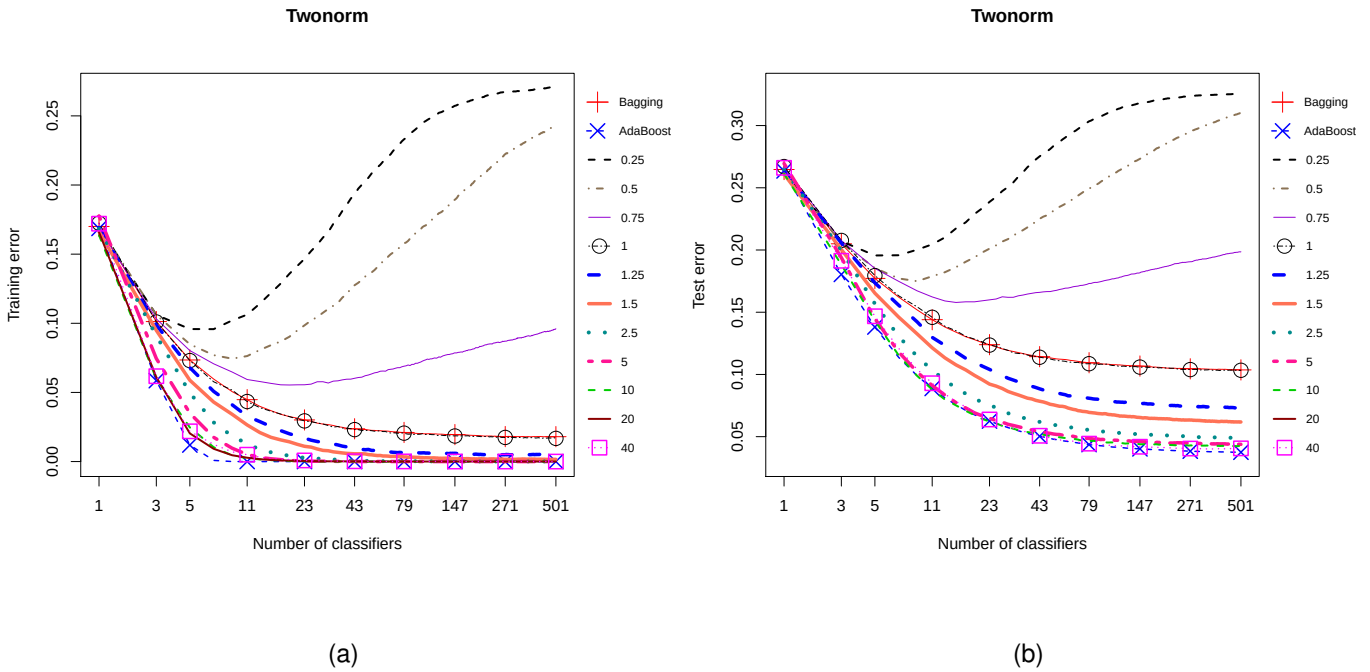

Fig. 4. Error rate as a function of number of classifiers in *Twonorm* for bagging, AdaBoost, and vote-boosting ensembles of pruned CART trees: (a) Training error (b) test error.

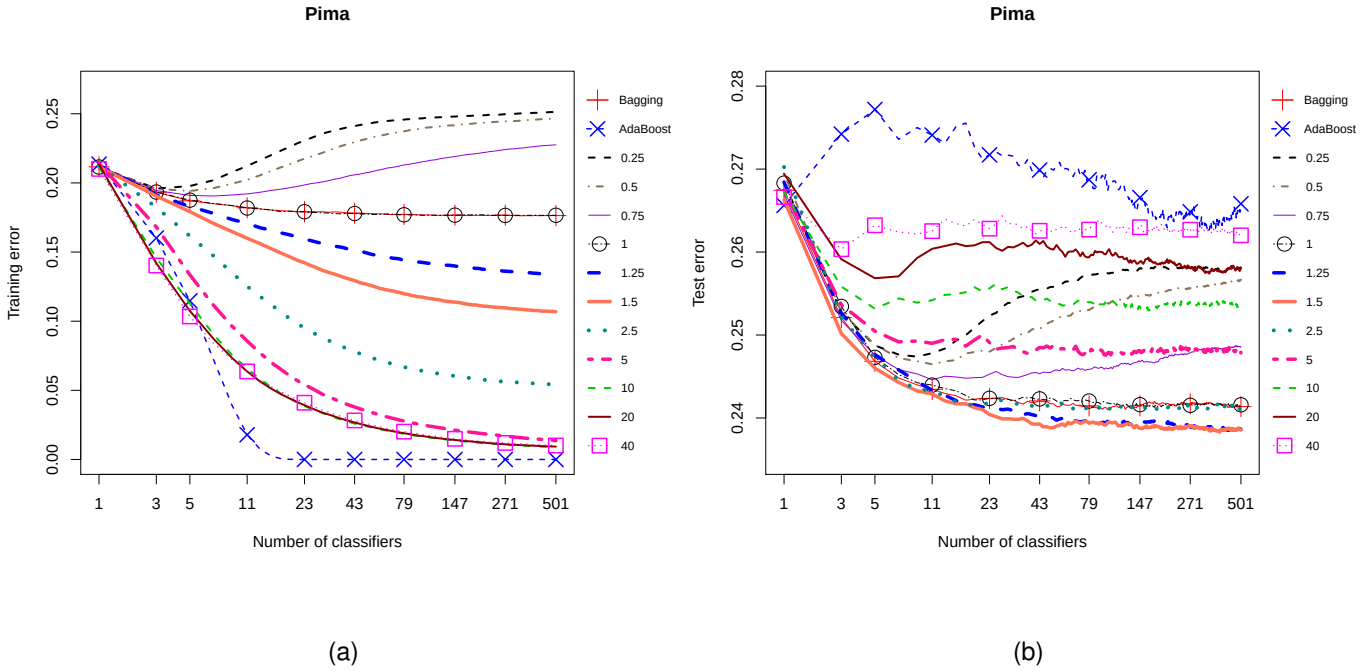

Fig. 5. Error rate as a function of number of classifiers in *Pima* for bagging, AdaBoost, and vote-boosting ensembles of pruned CART trees: (a) Training error (b) test error.

In this case, and if weighted resampling is used, vote-boosting is equivalent to bagging. This can be observed in Figures 4 and 5 as the evolution of the error for both method is almost equivalent.

When values of $a = b < 1.0$ are used, emphasis is made on instances on which the aggreement is highest. When regularized classifiers, such as pruned CART trees, are used as base learners, this type of emphasis is in general not effective. Nonetheless, as will be illustrated by the results presented in Section 4.2, focusing on these types of instances

can lead to improvements in the generalization capacity when the data are noisy and the ensemble is composed of classifiers that overfit.

Values of $a = b > 1.0$ correspond to emphasizing instances in which the ensemble prediction is uncertain. For Twonorm, the learning curves of vote-boosting using $a = b = 40.0$ and AdaBoost are similar. In this problem, the classification errors are mainly due to the fact that there is an overlap between the distributions of the two classes. Therefore, the incorrectly classified instances are close to the decision boundary, where the ensemble predictions are also more uncertain. Using a sharper beta distribution (see Fig. 1) gives more weights to these uncertain instances. In consequence, the resulting emphasis is similar to AdaBoost's. In *Pima*, which is a noisy problem, the learning curves of vote-boosting with large $a = b$ and Adaboost are different. But still, the closest performance to Adaboost is vote-boosting with large $a = b$. Finally, we observe that the optimal value for the shape parameter of the beta distribution in vote-boosting is problem dependent: Values of $a = b \approx 1.25 - 1.5$ perform well in *Pima*. The best performance in *Twonorm* requires using a large values of $a = b \approx 40$. For each problem, the optimal value can be determined using cross-validation.

4.2 Prediction accuracy

In this section we carry out an extensive comparison of the accuracy of vote-boosting ensembles with related ensemble methods: bagging, AdaBoost and random forest. In the empirical evaluation different classification tasks from the UCI repository [46] and other sources [47] are considered. In Table 1, the characteristics of the datasets used in our study are summarized. For each dataset the table displays the number of instances of the dataset, the number of instances used for training and for testing, and the number of attributes. All classification tasks are binary.

In each type of ensemble the base classifier that performs best is used: Pruned CART trees in AdaBoost, unpruned CART trees in bagging and random trees in vote-boosting and in random forest. We have used the adabag [48], ipred [49] and ranfomForest [50] packages in R for AdaBoost, bagging and random forest implementations, respectively. The ensembles considered are composed of 501 classifiers. This comparatively large ensemble size is needed in some problems to achieve convergence to the asymptotic error level [41]. Resampling with replacement is used to generate the bootstrap samples on which the individual classifiers are trained. In both vote-boosting and AdaBoost, weighted resampling is used to take into account the different emphasis on the training instances. In AdaBoost reweighting was considered as an alternative. However, as reported in the literature, similar or slightly better results are obtained when resampling instead of reweighting is used [12], [51].

In vote-boosting the symmetric beta distribution is used for emphasis. The shape parameter is determined as the one in $a = b \in \{0.25, 0.5, 0.75, 1, 1.25, 1.5, 2.5, 5, 10, 20, 40\}$ that minimizes the cross-validation error using 10 folds within the original training set. Since random trees are used as base learners, vote-boosting with uniform emphasis in all the training instances ($a = b = 1$) is equivalent to random forest. For values of the shape parameter above 1 the symmetric beta distribution has a single mode at $\pi = 0.5$, which implies that more emphasis is made on uncertain training instances. By contrast, when the shape parameter is smaller than 1 the focus is on training instances on which most classifiers agree.

The average test error rates for different classification problems of these ensemble methods are displayed in the columns of Table 2. The results reported are averages over 100 realizations of the training and test set partitions. For classification problems in which only a finite collection of instances are available $2/3$ of the data selected at random are used for training and the remaining $1/3$ for testing. For the synthetic problems instances are generated independently

TABLE 1
Characteristics of the classification problems analyzed and training / test partitions.

Dataset	Instances	Training	Test	Number of attributes
Australian	690	460	230	14
Breast W.	699	466	233	9
Blood	748	499	249	5
Boston	506	337	169	14
Chess	3196	2131	1065	36
German	1000	667	333	20
Heart	270	178	92	13
Hepatitis	155	104	51	19
Horse-colic	368	246	122	21
Ionosphere	351	234	117	34
Liver	345	230	115	6
Ozone	2536	1691	845	74
Parkinsons	197	132	65	24
Pima	768	512	256	8
Ringnorm	2300	300	2000	20
Sonar	208	139	69	60
Spambase	4601	3068	1533	58
Threenorm	2300	300	2000	20
Tic-tac-toe	958	639	319	9
Twonorm	2300	300	2000	20

TABLE 2
Average test error rates for bagging (unprunned CART trees), boosting (pruned CART trees), random forest and vote-boosting (random trees).

Dataset	Bagging	AdaBoost	Random forest	Vote-boosting	a = b (median)
Australian	13.3±2.0	13.6±1.9	**13.0±2.0**	13.2±1.9	1.2
Breast W.	4.6±1.2	3.5±0.9	**3.0±1.0***	3.2±1.0	1.0
Blood	26.2±1.9	25.3±2.2	24.6±2.1	**21.9±1.5***	0.5
Boston	14.1±2.6	**12.5±2.5***	13.0±2.3	13.1±2.3	2.5
Chess	0.6±0.3	**0.5±0.3***	1.6±0.4	0.6±0.2	10
German	24.4±1.9	24.6±2.1	**23.6±1.9***	24.2±1.8	2.5
Heart	19.7±3.7	20.2±3.3	17.1±3.3	**16.8±3.4**	0.5
Hepatitis	15.8±4.7	15.5±4.1	14.2±3.8	**13.8±3.7**	2.5
Horse-colic	**15.0±3.0***	21.2±3.3	**15.0±2.6***	**15.0±2.8***	1.5
Ionosphere	8.0±2.4	6.6±1.8	6.7±1.7	**6.4±1.7**	5
Liver	29.6±3.9	30.7±4.0	**27.9±3.5**	28.3±3.7	1.25
Ozone	8.6±0.6	**5.8±0.3***	6.0±0.2	5.9±0.2	7.5
Parkinsons	10.5±4.0	**7.2±3.4***	10.1±3.6	9.1±3.5	10
Pima	23.8±2.0	25.9±2.2	**23.4±2.0***	**23.4±1.9***	0.5
Ringnorm	8.9±1.8	**4.3±0.6***	6.0±1.1	4.4±0.8	20
Sonar	22.2±5.2	**15.1±4.3**	18.1±5.1	15.4±4.6	10
Spambase	6.0±0.5	4.4±0.5	5.0±0.5	**4.1±0.5***	20
Threenorm	19.0±1.8	16.8±0.8	16.7±0.9	**16.4±1.1***	5
Tic-tac-toe	2.1±1.0	**0.6±0.6***	2.2±0.9	1.2±0.7	10
Twonorm	6.4±1.5	4.0±0.6	3.9±0.5	**3.7±0.5***	5
win/draw/loss	**16/4/0**	**12/1/7**	**12/6/2**	—	

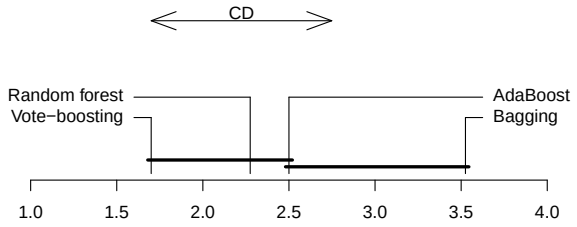

Fig. 6. Comparison of the average ranks of bagging, AdaBoost, random forest and vote-boosting using a Nemenyi test. Horizontal lines connect methods whose average ranks are not significantly different (p-value $<$ 0.05).

at random: 300 instances are used for training and 2000 for testing. In all cases stratified sampling is used to ensure that the class distributions in the test and training sets are similar to the whole set. In the last column, the median value for $a = b$ used in vote-boosting is reported. In Table 2 the best and the second best results for each classification problem are highlighted with bold face and underlined, respectively. In addition, the lowest test error rate is marked with an asterisk (*) if the improvement over to the second best is statistically significant. The significance of these differences is determined using a paired t-test at a significance level $\alpha = 0.05$. Finally, the number of statistically significant wins and losses when one compares the average accuracy of vote-boosting with each of the remaining ensembles are shown in the last row of the table. Draws correspond to differences of average accuracy that are not statistically significant.

Vote-boosting exhibits the best overall performance in terms of statistically significant wins. The tally is very favorable when one compares the average accuracy of vote-boosting with bagging: vote-boosting significantly outperforms bagging in 16 out of the 20 datasets and has an equivalent accuracy in *Australian*, *Chess*, *German* and *Horse-colic*. The number of significant wins in the comparison of vote-boosting and AdaBoost is 12 out of 20. In the datasets in which AdaBoost is the most accurate ensemble, with the exception of *Boston*, *Parkinsons* and *Tic-tac-toe*, the error rates of vote-boosting very close to those of AdaBoost. In these classification problems, the shape parameter of the beta distribution used as emphasis function in vote-boosting is fairly high (e.g. median values 7.5 in *Ozone*, 20 in *Ringnorm* and 10 in *Sonar*). This indicates a strong emphasis on uncertain examples, which has a similar effect as the emphasis on incorrectly classified instances that is characteristic of AdaBoost. On the other hand, in problems such as *Blood*, *Heart*, *Horse-colic*, *Liver* or *Pima*, which are difficult for AdaBoost, vote-boosting implements a milder or even negative emphasis, as evinced by the low values of the shape parameter. It is remarkable that in some cases values of a below 1 provide the best accuracy. In such cases, emphasis is made on instances on which the individual ensemble classifiers agree the most. The effectiveness of this type of emphasis, which is somewhat counterintuitive, is a consequence of the large intrinsic variability of random trees. In fact, this effect has not been found in ensembles of more stable base learners, such as pruned CART trees. Finally, vote-boosting is more accurate than random forest in 12 out of the 20 classification problems investigated. In the 2 classification tasks in which random forest outperforms vote-boosting (*Breast W.* and *German*) their differences, although statistically significant, are rather small. By contrast, in the problem cases where vote-boosting outperforms random forest, the improvements can be fairly large (e.g. *Blood*, *Chess*, *Ringnorm*, *Sonar*, *Tic-tac-toe*).

To provide an overall comparison of the accuracies of the different ensembles across the tested datasets we apply the procedure proposed by Demšar in [52]. In this comparison method one computes the average rank of the classifier in the 20 classification problems investigated. The rank of a classifier in a specific classification problem is determined by ordering the different methods according to their test errors. A lower rank corresponds to a smaller test error and, therefore, better accuracy. The average ranks of bagging and AdaBoost, random forest (RF) and vote-boosting are displayed in Fig. 4.2. In this diagram, the differences of average ranks between methods that are connected by a horizontal solid line are not statistically significant according to a Nemenyi test (p-value $<$ 0.05). From these results one concludes that bagging has the worst performance, in terms average rank. The differences with vote-boosting and random forest are statistically performance. In terms of this measure, vote-boosting obtains the best overall results. However, the differences with random forest and AdaBoost are not statistically significant.

From the values of the shape-parameter displayed in the last column of Table 2 it is apparent that the type of emphasis and its strength need to be adapted to the particular classification task at hand. To illustrate this point we have carried out a detailed analysis of the distribution of class votes for two classification task, in which the optimal emphasis strategies are very different.

In *Sonar* the optimal emphasis is to focus on training instances in which the disagreement rates are largest. The effects of this strategy is shown in Fig. 7. In this Fig., the distribution of votes in a vote-boosting ensemble of random trees built using a beta distribution with $a = b = 10.0$ can be analyzed. The plots in this figure display the evolution of histogram of the fraction of class votes for correctly (white bars) and incorrectly (red striped bars) classified instances in both the training set (upper row) and in the testing set (lower row), as a function of the size of the ensemble. The histograms shown correspond to vote-boosting ensembles of 1, 5, 11 and 501 random trees. In the training set the strong focus on uncertain instances (those for which the fraction of class votes is close to 0.5) leads to a markedly bimodal distribution, in which most predictions are by clear majority. The incorrectly classified instances disappear because the ensemble quickly achieves zero training error. The distribution of class votes in the test set is markedly different: It covers the whole interval, and has slightly higher frequencies for intermediate class vote frequencies, especially for the misclassified instances.

A very different picture is obtained in *Pima* (Fig. 8). In this classification task the selected shape parameter for the

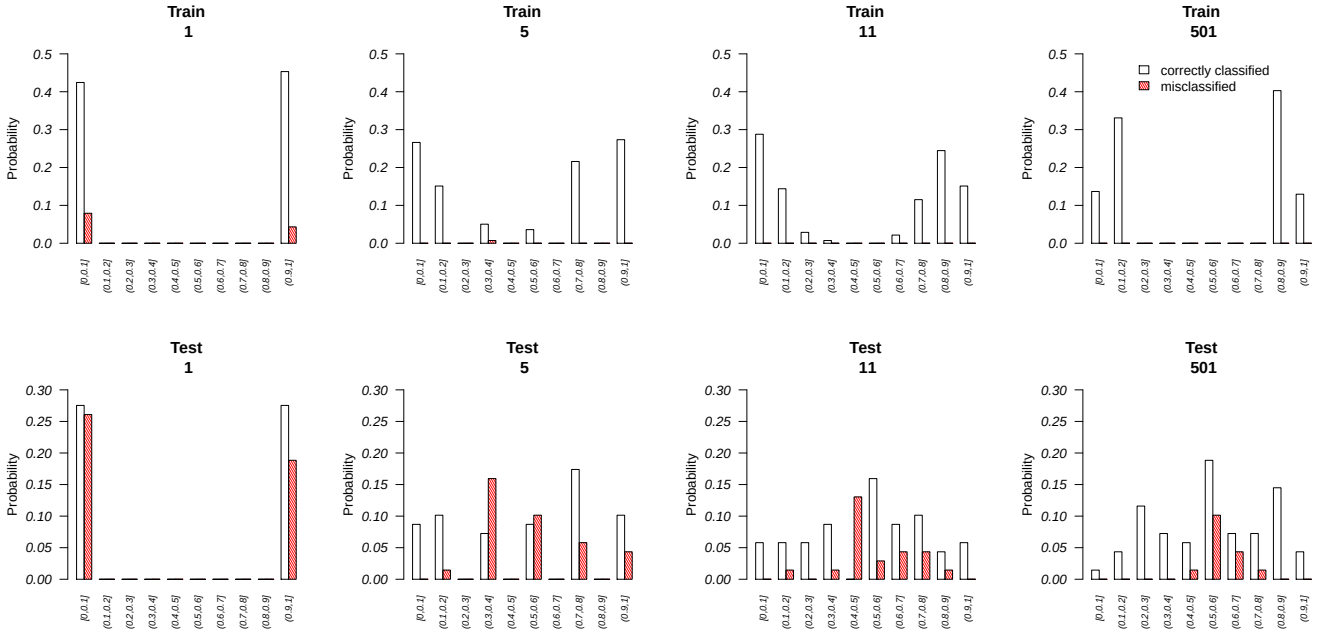

Fig. 7. Histograms of vote fractions for correctly (white) and incorrectly (red) classified training in *Sonar* for the training set (top row) and for the test set (bottom row)

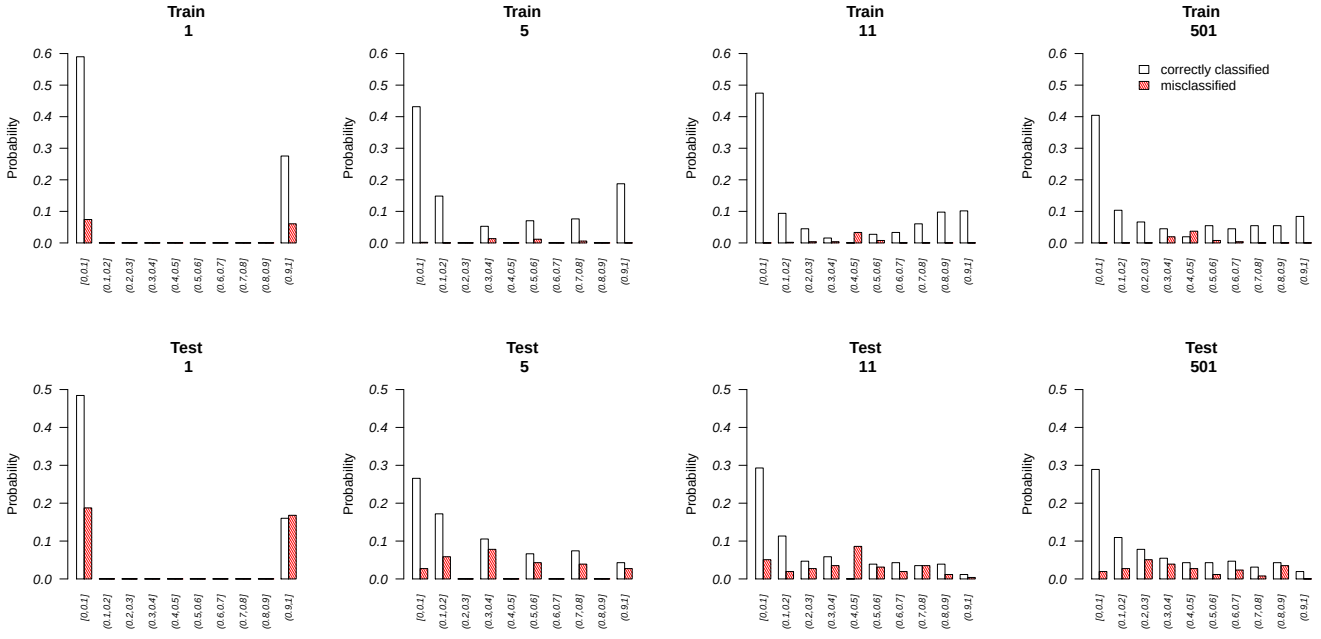

Fig. 8. Histograms of vote fractions for correctly (white) and incorrectly (red) classified training in *Pima* for the training set (top row) and for the test set (bottom row).

symmetric beta distribution is $a = b = 0.5$. In consequence, the optimal strategy is to avoid focusing on training instances in which the disagreement rates are highest. For the correctly classified instances the qualitative features in the training set and in the test set are similar. Misclassified instances in the training set appear mostly around 0.5. However, in the test set, they appear in the whole $[0, 1]$ interval, which is consistent with the observation that the problem seems to have levels of class-label noise.

5 Conclusions

Vote-boosting is a novel ensemble learning method in which individual classifiers are trained by progressively focusing on instances whose prediction is uncertain. This uncertainty is measured in terms of the disagreement rate among the

classifiers in the ensemble. For simple, regularized classifiers, such as decision stumps or pruned CART trees, vote-boosting interpolates between bagging and AdaBoost. When the level of class-label noise is small, prediction errors are more likely to occur near the classification boundary. Therefore, it is possible to build more accurate ensembles by focusing on uncertain instances. Since these instances are more likely to be misclassified, the emphasis given by vote-boosting is similar to AdaBoost's. For noisy classification problems, a softer emphasis on uncertain instances is generally preferable. In this case, the most accurate predictions are obtained by means of ensembles that are fairly similar to bagging.

When more brittle individual learners are used (e.g. random trees) a milder emphasis on uncertain instances is generally needed to achieve the best generalization performance. This is a consequence of the fact that some of the uncertainty in the predictions is due to the intrinsic variability of the base learners. For some noisy classification tasks one observes that it is in fact advantageous to reduce the relative weight of instances in which the disagreement rate is large.

Note that, since AdaBoost is based on emphasizing incorrectly classified instances, it cannot be used to improve the performance of base learners whose training error is small, such as random trees. By contrast, vote-boosting does not have this limitation and can be used to build effective ensembles composed of these types of classifiers.

ACKNOWLEDGMENTS

The research has been supported by the Spanish Ministerio de Economa y Competitividad (projects TIN2013-42351-P and TIN2015-70308-REDT), and the Comunidad de Madrid, project CASI-CAM-CM (S2013/ICE-2845).

REFERENCES

[1] T. G. Dietterich, "Ensemble Methods in Machine Learning," *Lecture Notes in Computer Science*, vol. 1857, pp. 1–15, 2000.

[2] L. Breiman, "Bagging Predictors," in *Machine Learning*, vol. 24, 1996, pp. 123–140.

[3] ——, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.

[4] Y. Liu and X. Yao, "Ensemble learning via negative correlation," *Neural Networks*, vol. 12, pp. 1399–1404, 1999.

[5] R. E. Schapire, "The strength of weak learnability," *Machine Learning*, vol. 5, pp. 197–227, 1990.

[6] Y. Freund and R. E. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," *Journal of Computer and System Sciences*, vol. 55, no. 1, pp. 119 – 139, 1997.

[7] J. Friedman, T. Hastie, and R. Tibshirani, "Additive logistic regression: a statistical view of boosting (With discussion and a rejoinder by the authors)," *The Annals of Statistics*, vol. 28, no. 2, pp. 337–407, Apr. 2000.

[8] R. E. Schapire and Y. Freund, *Boosting: Foundations and Algorithms*. The MIT Press, 2012.

[9] S. Tulyakov, S. Jaeger, V. Govindaraju, and D. Doermann, *Machine Learning in Document Analysis and Recognition*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, ch. Review of Classifier Combination Methods, pp. 361–386.

[10] G. Fumera and F. Roli, "A theoretical and experimental analysis of linear combiners for multiple classifier systems," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 27, no. 6, pp. 942–956, 2005.

[11] Y. Freund, R. Schapire, and N. Abe, "A short introduction to boosting," *Journal-Japanese Society For Artificial Intelligence*, vol. 14, no. 771-780, p. 1612, 1999.

[12] T. G. Dietterich, "An experimental comparison of three methods for constructing ensembles of decision trees: Bagging, boosting, and randomization," *Machine Learning*, vol. 40, no. 2, pp. 139–157, 2000.

[13] R. Maclin and D. W. Opitz, "Popular ensemble methods: An empirical study," *CoRR*, vol. abs/1106.0257, 2011.

[14] M. Fernández-Delgado, E. Cernadas, S. Barro, and D. Amorim, "Do we need hundreds of classifiers to solve real world classification problems?" *Journal of Machine Learning Research*, vol. 15, pp. 3133–3181, 2014.

[15] J. R. Quinlan, "Bagging, Boosting, and C4.5," in *In Proceedings of the Thirteenth National Conference on Artificial Intelligence*, 1996, pp. 725–730.

[16] E. Bauer and R. Kohavi, "An Empirical Comparison of Voting Classification Algorithms: Bagging, Boosting, and Variants," *Machine Learning*, vol. 36, no. 1, pp. 105–139, Jul. 1999.

[17] T. M. Khoshgoftaar, J. Van Hulse, and A. Napolitano, "Comparing Boosting and Bagging Techniques With Noisy and Imbalanced Data," *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, vol. 41, no. 3, pp. 552–568, May 2011.

[18] A. Loureiro, L. Torgo, and C. Soares, "Outlier detection using clustering methods: a data cleaning application," in *In Proceedings of the Data Mining for Business Workshop*, 2004.

[19] N. Abe, B. Zadrozny, and J. Langford, "Outlier detection by active learning," in *KDD '06: Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*. New York, NY, USA: ACM Press, 2006, pp. 767–772.

[20] C. Domingo and O. Watanabe, "MadaBoost: A modification of AdaBoost," in *Proceedings of the Thirteenth Annual Conference on Computational Learning Theory*, ser. COLT '00, 2000, pp. 180–189.

[21] W. Jiang, "Is regularization unnecessary for boosting?" in *Department of Statistics, Northwestern University*. Morgan Kaufmann, 2001, pp. 57–64.

[22] V. Gómez-Verdejo, J. Arenas-García, and A. R. Figueiras-Vidal, "A dynamically adjusted mixed emphasis method for building boosting ensembles," *IEEE Transactions on Neural Networks*, vol. 19, no. 1, pp. 3–17, 2008.

[23] P. K. Shivaswamy and T. Jebara, "Variance penalizing adaboost." in *NIPS*, J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. C. N. Pereira, and K. Q. Weinberger, Eds., 2011, pp. 1908–1916.

[24] G. Rätsch, T. Onoda, and K. R. Müller, "An improvement of adaboost to avoid overfitting," in *Proc. of the Int. Conf. on Neural Information Processing*. Citeseer, 1998.

[25] Y. Guo, P. L. Bartlett, A. J. Smola, R. C. Williamson, and J. Baxter, "Norm-based regularization of boosting," *Technical ReportAustralian National University*, 2002.

[26] Y. Sun, J. Li, and W. W. Hager, "Two new regularized adaboost algorithms," in *Proceedings of the 2004 International Conference on Machine Learning and Applications - ICMLA 2004, 16-18 December 2004, Louisville, KY, USA.*, 2004, pp. 41–48.

[27] Y. Sun, S. Todorovic, and J. Li, "Reducing the overfitting of adaboost by controlling its data distribution skewness," *IJPRAI*, vol. 20, no. 7, pp. 1093–1116, 2006.

[28] Y. Freund, "An adaptive version of the boost by majority algorithm," *Machine Learning*, vol. 43, no. 3, pp. 293–318, 2001.

[29] ——, "A more robust boosting algorithm," 2009. [Online]. Available: http://arxiv.org/abs/0905.2138v1

[30] S. Cheamanunkul, E. Ettinger, and Y. Freund, "Non-convex boosting overcomes random label noise," *CoRR*, vol. abs/1409.2905, 2014.

[31] R. A. McDonald, D. J. Hand, and I. A. Eckley, "An empirical comparison of three boosting algorithms on real data sets with artificial class noise," in *Multiple Classifier Systems, 4th International Workshop, MCS 2003, Guilford, UK, June 11-13, 2003, Proceedings*, 2003, pp. 35–44.

[32] V. Gómez-Verdejo, M. Ortega-Moral, J. Arenas-García, and A. R. Figueiras-Vidal, "Boosting by weighting critical and erroneous samples," *Neurocomputing*, vol. 69, no. 7-9, pp. 679–685, 2006.

[33] V. Gómez-Verdejo, J. Arenas-García, and A. R. Figueiras-Vidal, "A dynamically adjusted mixed emphasis method for building boosting ensembles," *IEEE Trans. Neural Networks*, vol. 19, no. 1, pp. 3–17, 2008.

[34] A. Ahachad, A. Omari, and A. R. Figueiras-Vidal, "Neighborhood guided smoothed emphasis for real adaboost ensembles," *Neural Processing Letters*, vol. 42, no. 1, pp. 155–165, 2015.

[35] P. M. Long and R. A. Servedio, "Random classification noise defeats all convex potential boosters," *Machine Learning*, vol. 78, no. 3, pp. 287–304, 2010.
[36] G. Martinez-Munoz and A. Suarez, "Switching class labels to generate classification ensembles," *Pattern Recognition*, vol. 38, no. 10, pp. 1483–1494, 2005.
[37] L. K. Hansen and P. Salamon, "Neural network ensembles," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 12, no. 10, pp. 993–1001, Oct. 1990.
[38] J. J. Rodríguez, L. I. Kuncheva, and C. J. Alonso, "Rotation forest: A new classifier ensemble method." *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 10, pp. 1619–1630, 2006.
[39] P. Geurts, D. Ernst, and L. Wehenkel, "Extremely randomized trees," *Machine Learning*, vol. 63, no. 1, pp. 3–42, 2006.
[40] D. Hernández-Lobato, G. Martinez-Muoz, and A. Suarez, "Statistical instance-based pruning in ensembles of independent classifiers," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 31, no. 2, pp. 364–369, 2009.
[41] D. Hernández-Lobato, G. Martínez-Muñoz, and A. Suárez, "Inference on the prediction of ensembles of infinite size," *Pattern Recognition*, vol. 44, no. 7, pp. 1426 – 1434, 2011.
[42] D. Hernández-Lobato, G. Martínez-Muñoz, and A. Suárez, "How large should ensembles of classifiers be?" *Pattern Recognition*, vol. 46, no. 5, pp. 1323–1336, May 2013.
[43] L. Mason, J. Baxter, P. Bartlett, and M. Frean, "Boosting algorithms as gradient descent," in *In Advances in Neural Information Processing Systems 12*. MIT Press, 2000, pp. 512–518.
[44] L. Breiman, "Bias, Variance and Arcing Classifiers," UC Berkeley, Tech. Rep. 460, 1996.
[45] B. Frénay and M. Verleysen, "Classification in the presence of label noise: A survey," *IEEE Trans. Neural Netw. Learning Syst.*, vol. 25, no. 5, pp. 845–869, 2014.
[46] K. Bache and M. Lichman, "UCI machine learning repository," 2013. [Online]. Available: http://archive.ics.uci.edu/ml
[47] L. Breiman, "Arcing classifiers," *Annals of Statistics*, vol. 26, pp. 801–823, 1998.
[48] E. Alfaro, M. Gámez, and N. García, "adabag: An R package for classification with boosting and bagging," *Journal of Statistical Software*, vol. 54, no. 2, pp. 1–35, 2013.
[49] A. Peters and T. Hothorn, "ipred: improved predictors. r package version 0.8-8," 2009.
[50] A. Liaw and M. Wiener, "Classification and regression by randomforest," *R News*, vol. 2, no. 3, pp. 18–22, 2002.
[51] C. Seiffert, T. M. Khoshgoftaar, J. V. Hulse, and A. Napolitano, "Resampling or Reweighting: A Comparison of Boosting Implementations," in *2008 20th IEEE International Conference on Tools with Artificial Intelligence*. IEEE, Nov. 2008, pp. 445–451.
[52] J. Demšar, "Statistical comparisons of classifiers over multiple data sets," *J. Mach. Learn. Res.*, vol. 7, pp. 1–30, Dec. 2006.

Maryam Sabzevari Maryam Sabzevari received her B.Sc. (2008) and M.Sc. (2010) in computer science from Azad university of Mashhad, Iran. Currently she is conducting her Ph.D. in Computer Science in Universidad Autnoma de Madrid (UAM), Madrid, Spain. Her interest include machine learning, pattern recognition, decision trees and ensemble learning.

Gonzalo Martínez-Muñoz received the university degree in Physics (1995) and Ph.D. degree in Computer Science (2006) from the Universidad Autónoma de Madrid (UAM). From 1996 to 2002 he worked in industry. Until 2008 he was an interim assistant professor in the Computer Science Department of the UAM. During 2008/2009, he worked as a Fulbright postdoc researcher at Oregon State University in the group of Professor Thomas G. Dietterich. He is currently a professor at Computer Science Department at UAM. His research interests include machine learning, computer vision, pattern recognition, neural networks, decision trees, and ensemble learning.

Alberto Suárez received the degree of Licenciado in Chemistry from the Universidad Autónoma de Madrid, Spain, in 1988, and the Ph.D. degree in Physical Chemistry from the Massachusetts Institute of Technology (MIT), Cambridge, MA, in 1993. After holding postdoctoral positions at Stanford University (USA), at the Université Libre de Bruxelles (Belgium), as a research fellow financed by the European Commission within the program Training and Mobility of Researchers, and at the Katholieke Universiteit Leuven (Belgium), he is currently a professor in the Computer Science Department of the Universidad Autónoma de Madrid (Spain). He has also held appointments as Senior Visiting Scientist at the International Computer Science Institute (Berkeley, CA) and at MIT (Cambridge, MA). He has worked on relaxation theory in condensed media, stochastic and thermodynamic theories of nonequilibrium systems, lattice-gas automata, and automatic induction from data. His current research interests include machine learning, quantitative and computational finance, time series analysis and information processing in the presence of noise. He is a member of IEEE.